

Claims

1. A computer-implemented method for shaping the output of cells on an output path of a data transmitting device, the data transmitting device being configured for regulating the cells from a plurality of input paths to the output path to a network, comprising:

5 sorting a plurality of queues, each queue including a plurality of cells associated with a communication device, wherein the plurality of queues are arranged according to a weight and a data rate associated with each plurality of cells resulting in a plurality of sorted queues of queues;

10 regulating an aggregate output of cells from each sorted queue of queues based upon the data rates of the queues of the each sorted queue of queues; and

scheduling the regulated output of the aggregate output of cells from each sorted queue of queues based upon the weights of the each sorted queue of queues, such that the scheduled output is coupled to the output path;

15 wherein the scheduled output conforms to a plurality of characteristics of the network, such that the network is efficiently used to carry the cells from the plurality of input paths to a plurality of communication devices.

2. A method as claimed in claim 1, wherein said plurality of input paths represent virtual
20 connections.

B
cancel

3. A method as claimed in claim 1, wherein said plurality of input paths represent input ports

5 4. A computer-implemented method for shaping the output of cells on an output path of a data transmitting device, the data transmitting device being configured for switching the cells from a plurality of input paths to the output path, comprising:

receiving a plurality of cells through the plurality of input paths, the cells arriving from a plurality of communication devices, each cell having a weight and an associated data rate;

10 sorting each cell of the plurality of cells into a plurality of queues of cells, each queue of the plurality of queues corresponding to one of the plurality of communication devices, the cells of each queue having the same weight and data rate;

15 sorting each queue of the plurality of queues into a first plurality of queues of queues according to the weight of each queue, such that each queue of queues of the first plurality of queues of queues includes queues having the same weight;

sorting each queue of each queue of queues of the first plurality of queues of queues into a second plurality of queues of queues according to the data rate of each queue;

20 shaping an output of cells from each queue of each queue of queues of the second plurality of queues of queues, shaping being implemented by a plurality of traffic shapers, each traffic shaper corresponding to a queue of queues of the second plurality of queue of queues

whereby the corresponding queue of queues have the same weight, which determines a weight of the each traffic shaper; and

scheduling an output of each traffic shaper, scheduling being implemented by a scheduler, such that the scheduled output is coupled to the output path.

5

5. A method as claimed in claim 4, wherein each queue of queues of the second plurality of queues of queues is comprised of a first subset of queues having approximately the same data rate.

10

6. A method as claimed in claim 5, wherein each queue of queues of the second plurality of queues of queues is further comprised of a second subset of queues, such that the aggregate data rate of the second set of queues is approximately equal to the individual data rate of each queue of the first set of queues.

15

7. A method as claimed in claim 6, wherein in the step of shaping, each traffic shaper receives cells from the corresponding queue of queues of the second plurality of queues of queues, such that each traffic shaper receives cells from each queue of the first subset of queues of the corresponding queue of queues at approximately the same rate, and aggregately receives cells from the second subset of queues at approximately the same rate as each queue of the first

20

subset of queues.

8. A method as claimed in claim 7, wherein in the step of shaping, each traffic shaper receives cells from the corresponding queue of queues of the second plurality of queues of queues per a predetermined time interval associated with the data rates of the corresponding queue of queues, such that each traffic shaper receives one cell from each of queue of the first subset of queues of the corresponding queue of queues and receives a cell from a scheduled queue of the second subset of queues per predetermined time interval, such that the each queue of the second subset of queues are sequentially alternated as the scheduled queue per predetermined time interval.

9. A method as claimed in claim 4, wherein in the step of shaping, each traffic shaper allows each queue of each queue of queues of the second plurality of queues of queues to output data based upon the weight and data rate of the each queue of queues.

10. A method as claimed in claim 9, wherein each traffic shaper allows each queue of each queue of queues of the second plurality of queues of queues to output data when the each traffic shaper triggers, the triggering of the each traffic shaper including:

providing a threshold level for each traffic shaper;

providing a credit for each traffic shaper, the credit being determined by the data rate of the queue of queues of the second plurality of queues of queues corresponding to the each traffic shaper;

providing a bucket level for each traffic shaper;

incrementing the bucket level by the credit at a predetermined time interval associated with the data rate of the queue of queues of the second plurality of queues of queues corresponding to the each traffic shaper;

5 triggering the each traffic shaper when the bucket level is greater than the threshold level;

decrementing the bucket level of the each traffic shaper when the each traffic shaper allows each queue of each queue of queues of the second plurality of queues of queues to output data.

11. A method as claimed in claim 4, the scheduler including,

10 a plurality of round robin schedulers, wherein each round robin scheduler schedules a subset of the plurality of traffic shapers to output cells, whereby each traffic shaper of the subset of traffic shapers are of the same weight, which determines a weight of the round robin scheduler, wherein the each round robin scheduler schedules output from the subset of traffic shapers by the order in which each traffic shaper of the subset of traffic shapers triggers; and

15 a priority scheduler, wherein the priority scheduler schedules the plurality of round robin schedulers to output cells, wherein the priority scheduler schedules the output from the plurality of round robin schedulers by the weights of the plurality of round robin schedulers.

12. A method as claimed in claim 10, the scheduler including

a plurality of round robin schedulers, wherein each round robin scheduler schedules a subset of the plurality of traffic shapers to output cells, whereby each traffic shaper of the subset of traffic shapers are of the same weight, which determines a weight of the round robin scheduler, wherein the each round robin scheduler schedules output from the subset of traffic shapers by the order in which each traffic shaper of the subset of traffic shapers triggers; and

a priority scheduler, wherein the priority scheduler schedules the plurality of round robin schedulers to output cells, wherein the priority scheduler schedules the output from the plurality of round robin schedulers by the weights of the plurality of round robin schedulers.

13. A communication device, wherein the communication device shapes traffic according to the method recited in claim 4.

14. An interface card, wherein the interface card shapes traffic according to the method recited in claim 4.

15. A network card, wherein the network card shapes traffic according to the method recited in claim 4.

16. A network switch as claimed in claim 15, wherein the network switch is an asynchronous transfer mode switch.

17. A computer-implemented method for shaping the output of cells on an output path of a data switch, the data switch being configured for switching the cells from a plurality of input paths to the output path, comprising:

5 receiving a plurality of cells through the plurality of input path, the cells arriving from a plurality of communication devices, each cells having a weight and an associated data rate;

sorting each cell of the plurality of cells into a plurality of queues of cells, each queue of the plurality of queues corresponding to one of the plurality of communication devices, the cells of each queue having the same weight and data rate;

10 sorting each queue of the plurality of queues into a first plurality of queues of queues according to the weight of each queues, such that each queue of queues of the first plurality of queues of queues includes queues having the same weight;

sorting each queue of each queue of queues of the first plurality of queues of queues into a second plurality of queues of queues according to the data rate of each queue, wherein each
15 queue of queues of the second plurality of queues of queues is comprised of a first subset of queues having approximately the same data rate and a second subset of queues, such that the aggregate data rate of the second set of queues is approximately equal to the individual data rate of each queue of the first set of queues;

shaping an output of cells from each queue of each queue of queues of the second
20 plurality of queues of queues, shaping being implemented by a plurality of traffic shapers, each

traffic shaper corresponding to a queue of queues of the second plurality of queue of queues, shaping including,

each traffic shaper receiving cells from the corresponding queue of queues of the second plurality of queues of queues, such that each traffic shaper receives cells from each queue of the first subset of queues of the corresponding queue of queues at approximately the same rate, and aggregately receiving cells from the second subset of queues at approximately the same rate as each queue of the first subset of queues, and

each traffic shaper allowing each queue of each queue of queues of the second plurality of queues of queues to output data based upon the weight and data rate of the each queue of queues, whereby the corresponding queue of queues have the same weight, which determines a weight of the each traffic shaper; and

scheduling an output of each traffic shaper, scheduling being implemented by a scheduler, the scheduler including a plurality of round robin schedulers, wherein each round robin scheduler schedules a subset of the plurality of traffic shapers to output cells, whereby each traffic shaper of the subset of traffic shapers are of the same weight, which determines a weight of the round robin scheduler, wherein the each round robin scheduler schedules output from the subset of traffic shapers by the order in which each traffic shaper of the subset of traffic shapers triggers, and a priority scheduler, wherein the priority scheduler schedules the plurality of round robin schedulers to output cells, wherein the priority scheduler schedules the output from the plurality of round robin schedulers by the weights of the plurality of round robin schedulers, such that the scheduled output is coupled to the output path,

18. A method as claimed in claim 17, wherein in the substep of each traffic shaper receiving cells from the corresponding queue of queues of the second plurality of queues of queues, each traffic shaper receives cells from the corresponding queue of queues per a predetermined time interval associated with the data rates of the corresponding queue of queues, such that each traffic shaper receives one cell from each of queue of the first subset of queues of the corresponding queue of queues and receiving a cell from a scheduled queue of the second subset of queues per predetermined time interval, such that the each queue of the second subset of queues are sequentially alternated as the scheduled queue per predetermined time interval

19. A method as claimed in claim 17, wherein in the substep of each traffic shaper allowing each queue of queues of the second plurality of queue of queues to output data, each traffic shaper allows each queue of each queue of queues of the second plurality of queues of queues to output data when the each traffic shaper triggers, the triggering of the each traffic shaper including:

providing a threshold level for each traffic shaper;

providing a credit for each traffic shaper, the credit being determined by the data rate of the queue of queues of the second plurality of queues of queues corresponding to the each traffic shaper;

providing a bucket level for each traffic shaper;

incrementing the bucket level by the credit at a predetermined time interval associated with the data rate of the queue of queues of the second plurality of queues of queues corresponding to the each traffic shaper;

triggering the each traffic shaper when the bucket level is greater than the threshold level;

5 decrementing the bucket level of the each traffic shaper when the each traffic shaper allows each queue of each queue of queues of the second plurality of queues of queues to output data.

10 20. A computer-implemented method for scheduling the output of cells on an output path of a data switch, the data switch being configured for switching the cells from a plurality of input paths to the output path, comprising:

providing a plurality of queues, each queue of said plurality of queues having a weight and an associated data rate, respective ones of said plurality of input paths being coupled to respective ones of said plurality of queues;

15 providing a plurality of queues of queues, the plurality of queues being coupled to the plurality of queues of queues, such that the queues of the plurality of queues having a similar weight being coupled to a same queue of queues of the plurality of queue of queues;

providing a plurality of traffic shapers, each of the plurality of queues of queues being input to an associated traffic shaper;

20 providing a scheduler, the plurality of traffic shapers being input to the scheduler, the scheduler being coupled to said output path.

21. A network switch for receiving a plurality of cells from a first plurality of communication devices from a plurality of input paths and outputting the cells to an output path to a second plurality of communication devices, wherein a plurality of virtual circuits are created between the first and second plurality of communication devices, each virtual circuit having a weight and a data rate, comprising:

a virtual circuit manager, the virtual circuit manager sorting the plurality of cells into a plurality of queues, each queue including cells corresponding to a one of the plurality of virtual circuits, arranging the plurality of queues into a first plurality of queues of queues, each queue of queues including queues having cells corresponding to virtual circuits having the same weight, and further arranging each of the first plurality of queues of queues into a second plurality of queues of queues, each queue of queues of the second plurality of queues of queues having cells corresponding to a first subset of virtual circuits having a first data rate and second subset of virtual circuits having an aggregate data rate approximately equal to the first data rate;

a plurality of traffic shaper link lists, each traffic shaper link list organizing a corresponding queue of queues of the second plurality of queues of queues, wherein the each traffic shaper includes

a non-aggregate linklist, the aggregate linklist keeping track of a sequence of a first subset of the second plurality of queues of queues corresponding to the first subset of virtual circuits having the first data rate,

a non-aggregate linklist, the non-aggregate linklist keeping track of a sequence of a second subset of the second plurality of queues of queues corresponding to the second subset of virtual circuits having an aggregate data rate approximately equal to the first data rate, and

a link table, the link table keeping track of a beginning and an end of the first and
5 second sequence, wherein each traffic shaper link list also schedules the queues of the corresponding queue of queues of the second plurality of queues of queues to output cells;

a plurality of traffic shapers, each of the plurality of traffic shapers corresponding to a one of the plurality of traffic shaper link lists, each traffic shaper controlling the scheduled output of cells from the corresponding queue of queues of the second plurality of queues of
10 queues, wherein the traffic shaper allows the corresponding queue of queues to output cells when the traffic shaper is triggered and the traffic shaper is given permission; and

a scheduler, the scheduler giving permission to the plurality of traffic shapers, wherein permission is given to a plurality of subsets of the plurality of traffic shapers, each subset of the plurality of traffic shapers corresponding to a queue of queues of the first plurality of queues of
15 queues, according to the weight of the cells within the corresponding queue of queues of the first plurality of queues of queues;

such that the output of cells is coupled to the output port and the sequence of the output of cells is shaped to take advantage of a plurality of characteristics of each of the virtual circuits.